

01-03-00

A

## FISH &amp; RICHARDSON P.C.

4225 Executive Square  
Suite 1400  
La Jolla, California  
92037

Telephone  
858 678-5070

Facsimile  
858 678-5099

Web Site  
www.fr.com

PTO  
09/475614

12/30/99

December 30, 1999

Attorney Docket No.: 10559/137001/P7876

**Box Patent Application**

Assistant Commissioner for Patents  
Washington, DC 20231

Presented for filing is a new original patent application of:

Applicant: GILBERT WOLRICH, DEBRA BERNSTEIN, MATTHEW J.  
ADILETTA AND DONALD F. HOOPER

Title: METHOD AND APPARATUS FOR CONTROL OF RECEIVE DATA

Enclosed are the following papers, including those required to receive a filing date  
under 37 CFR 1.53(b):

	Pages
Specification	35
Claims	8
Abstract	1
Declaration	[To be Filed at a Later Date]
Drawing(s)	13

Enclosures:

— Postcard.

Basic filing fee	\$0
Total claims in excess of 20 times \$18	\$0
Independent claims in excess of 3 times \$78	\$0
Fee for multiple dependent claims	\$0
Total filing fee:	\$0

## CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL528184943US

I hereby certify that this correspondence is being deposited with the  
United States Postal Service as Express Mail Post Office to Addressee  
with sufficient postage on the date indicated below and is addressed to  
the Assistant Commissioner for Patents, Washington, D.C. 20231.

Date of Deposit December 30, 1999

Signature

Deborah L. Dean

Typed or Printed Name of Person Signing Certificate

12/30/99  
jc712 U.S. PTO  
Frederick P. Fish  
1855-1930  
W.K. Richardson  
1859-1951

BOSTON  
DELAWARE  
NEW YORK  
SILICON VALLEY  
SOUTHERN CALIFORNIA  
TWIN CITIES  
WASHINGTON, DC

FISH & RICHARDSON P.C.

Assistant Commissioner for Patents  
December 30, 1999  
Page 2

No filing fee is being paid at this time. Please apply any other required fees, **EXCEPT FOR THE FILING FEE**, to deposit account 06-1050, referencing the attorney document number shown above. A duplicate copy of this transmittal letter is attached.

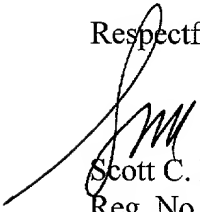
If this application is found to be incomplete, or if a telephone conference would otherwise be helpful, please call the undersigned at (858) 678-5070.

Kindly acknowledge receipt of this application by returning the enclosed postcard.

Please send all correspondence to:

SCOTT C. HARRIS  
Fish & Richardson P.C.  
4225 Executive Square, Suite 1400  
La Jolla, CA 92037

Respectfully submitted,



Scott C. Harris  
Reg. No. 32,030  
Enclosures  
SCH/rpi  
10012888.doc

APPLICATION  
FOR  
UNITED STATES LETTERS PATENT

TITLE: METHOD AND APPARATUS FOR CONTROL OF  
RECEIVE DATA

APPLICANT: GILBERT WOLRICH, DEBRA BERNSTEIN, MATTHEW J.  
ADILETTA AND DONALD F. HOOPER

CERTIFICATE OF MAILING BY EXPRESS MAIL

Express Mail Label No. EL528184943US

I hereby certify that this correspondence is being deposited with the United States Postal Service as Express Mail Post Office to Addressee with sufficient postage on the date indicated below and is addressed to the Assistant Commissioner for Patents, Washington, D.C. 20231.

December 30, 1999  
Date of Deposit

Signature

Deborah L. Dean

Typed or Printed Name of Person Signing Certificate

METHOD AND APPARATUS FOR CONTROL OF RECEIVE DATA

Background of the Invention

5 The invention relates generally to network data processing.

10 Networking products such as routers require high speed components for packet data movement, i.e., collecting packet data from incoming network device ports and queuing the packet data for transfer to appropriate forwarding device ports. They also require high-speed special controllers for processing the packet data, that is, parsing the data and making forwarding decisions. Because the implementation of these high-speed functions usually involves the development of ASIC or custom devices, such  
15 networking products are of limited flexibility. For example, each controller is assigned to service network packets from for one or more given ports on a permanent basis.

20 Summary of the Invention

In one aspect of the invention, receiving data from

a network includes issuing a receive request directing the transfer of data from one of the plurality of device ports to a buffer memory and specifying a thread from among a plurality of processing threads to process the data.

5

Brief Description of the Drawings

Other features and advantages of the invention will be apparent from the following description taken together with the drawings in which:

FIG. 1 is a block diagram of a communication system employing a hardware-based multi-threaded processor;

FIG. 2 is a block diagram of a microengine employed in the hardware-based multi-threaded processor of FIG. 1;

FIG. 3 is an illustration of an exemplary thread task assignment;

FIG. 4 is a block diagram of an I/O bus interface shown in FIG. 1;

FIG. 5 is a detailed diagram of a bus interface unit employed by the I/O bus interface of FIG. 4;

FIGS. 6A-6F are illustrations of various bus configuration control and status registers (CSRs);

FIG. 7A is a detailed diagram illustrating the interconnection between a plurality of 10/100 Ethernet ("slow") ports and the bus interface unit;

FIG. 7B is a detailed diagram illustrating the interconnection between two Gigabit Ethernet ("fast") ports and the bus interface unit;

FIGS. 8A-8C are illustrations of the formats of the RCV\_\_RDY\_CTL, RCV\_RDY\_HI and RCV\_RDY\_LO CSR registers, respectively;

FIG. 9 is a depiction of the receive threads and their interaction with the I/O bus interface during a receive process;

FIGS. 10A and 10B are illustrations of the format of the RCV\_REQ FIFO and the RCV\_CTL FIFO, respectively; and

FIG. 11 is an illustration of the thread done registers.

### Detailed Description

Referring to FIG. 1, a communication system 10 includes a parallel, hardware-based multi-threaded processor 12. The hardware based multi-threaded processor 12 is

coupled to a first peripheral bus (shown as a PCI bus) 14, a second peripheral bus referred to as an I/O bus 16 and a memory system 18. The system 10 is especially useful for tasks that can be broken into parallel subtasks or functions. The hardware-based multi-threaded processor 12 includes multiple microengines 22, each with multiple hardware controlled program threads that can be simultaneously active and independently work on a task. In the embodiment shown, there are six microengines 22a-22f and each of the six microengines is capable of processing four program threads, as will be described more fully below.

The hardware-based multi-threaded processor 12 also includes a processor 23 that assists in loading microcode control for other resources of the hardware-based multi-threaded processor 12 and performs other general purpose computer type functions such as handling protocols, exceptions, extra support for packet processing where the microengines pass the packets off for more detailed processing. In one embodiment, the processor 23 is a StrongARM (ARM is a trademark of ARM Limited, United Kingdom) core based architecture. The processor (or core) 23 has an operating system through which the processor 23

can call functions to operate on the microengines 22a-22f.  
The processor 23 can use any supported operating system,  
preferably real-time operating system. For the core  
processor implemented as a StrongARM architecture, operating  
systems such as MicrosoftNT real-time, VXWorks and :CUS, a  
freeware operating system available over the Internet, can  
be used.

The six microengines 22a-22f each operate with  
shared resources including the memory system 18, a PCI bus  
interface 24 and an I/O bus interface 28. The PCI bus  
interface provides an interface to the PCI bus 14. The I/O  
bus interface 28 is responsible for controlling and  
interfacing the processor 12 to the I/O bus 16. The memory  
system 18 includes a Synchronous Dynamic Random Access  
Memory (SDRAM) 18a, which is accessed via an SDRAM  
controller 26a, a Static Random Access Memory (SRAM) 18b,  
which is accessed using an SRAM controller 26b, and a  
nonvolatile memory (shown as a FlashROM) 18c that is used  
for boot operations. The SDRAM 16a and SDRAM controller 26a  
are typically used for processing large volumes of data,  
e.g., processing of payloads from network packets. The SRAM  
18b and SRAM controller 26b are used in a networking



implementation for low latency, fast access tasks, e.g.,  
accessing look-up tables, memory for the processor 23, and  
so forth. The microengines 22a-22f can execute memory  
reference instructions to either the SDRAM controller 26a or  
the SRAM controller 18b.

The hardware-based multi-threaded processor 12  
interfaces to network devices such as a media access  
controller device, including a "slow" device 30 (e.g.,  
10/100BaseT Ethernet MAC) and/or a "fast" device 31, such as  
Gigabit Ethernet MAC, ATM device or the like, over the I/O  
Bus 16. In the embodiment shown, the slow device 30 is an  
10/100 BaseT Octal MAC device and thus includes 8 slow ports  
32a-32h, and the fast device is a Dual Gigabit MAC device  
having two fast ports 33a, 33b. Each of the network devices  
attached to the I/O Bus 16 can include a plurality of ports  
to be serviced by the processor 12. Other devices, such as  
a host computer (not shown), that may be coupled to the PCI  
bus 14 are also serviced by the processor 12. In general,  
as a network processor, the processor 12 can interface to  
any type of communication device or interface that  
receives/sends large amounts of data. The processor 12  
functioning as a network processor could receive units of

packet data from the devices 30, 31 and process those units  
of packet data in a parallel manner, as will be described.  
The unit of packet data could include an entire network  
packet (e.g., Ethernet packet) or a portion of such a  
packet.

Each of the functional units of the processor 12 are  
coupled to one or more internal buses. The internal buses  
include an internal core bus 34 (labeled "AMBA") for  
coupling the processor 23 to the memory controllers 26a, 26b  
and to an AMBA translator 36. The processor 12 also  
includes a private bus 38 that couples the microengines 22a-  
22f to the SRAM controller 26b, AMBA translator 36 and the  
Fbus interface 28. A memory bus 40 couples the memory  
controllers 26a, 26b to the bus interfaces 24, 28 and the  
memory system 18.

Referring to FIG. 3, an exemplary one of the  
microengines 22a-22f is shown. The microengine 22a includes  
a control store 70 for storing a microprogram. The  
microprogram is loadable by the central processor 20. The  
microengine 70 also includes control logic 72. The control  
logic 72 includes an instruction decoder 73 and program  
counter units 72a-72d. The four program counters are

maintained in hardware. The microengine 22a also includes context event switching logic 74. The context event switching logic 74 receives messages (e.g., SEQ\_#\_EVENT\_RESPONSE; FBI\_EVENT\_RESPONSE; SRAM\_EVENT\_RESPONSE; SDRAM\_EVENT\_RESPONSE; and AMBA\_EVENT\_RESPONSE) from each one of the share resources, e.g., SRAM 26b, SDRAM 26a, or processor core 20, control and status registers, and so forth. These messages provides information on whether a requested function has completed. Based on whether or not the function requested by a thread has completed and signaled completion, the thread needs to wait for that complete signal, and if the thread is enabled to operate, then the thread is place on an available thread list (not shown). As earlier mentioned, the microengine 22a can have a maximum of 4 threads of execution available.

In addition to event signals that are local to an executing thread, the microengine employs signaling states that are global. With signaling states, an executing thread can broadcast a signal state to all microengines 22. Any and all threads in the microengines can branch on these signaling states. These signaling states can be used to determine availability of a resource or whether a resource

is due for servicing.

5 The context event logic 74 has arbitration for the  
four threads. In one embodiment, the arbitration is a round  
robin mechanism. However, other arbitration techniques,  
such as priority queuing or weighted fair queuing, could be  
used. The microengine 22a also includes an execution box  
(EBOX) data path 76 that includes an arithmetic logic unit  
(ALU) 76a and a general purpose register (GPR) set 76b. The  
ALU 76a performs arithmetic and logical functions as well as  
10 shift functions.

15 The microengine 22a further includes a write  
transfer register file 78 and a read transfer register file  
80. The write transfer register file 78 stores data to be  
written to a resource. The read transfer register file 80  
is for storing return data from a resource. Subsequent to  
or concurrent with the data arrival, an event signal from  
the respective shared resource, e.g., memory controllers  
26a, 26b, or core 23, will be provided to the context event  
arbiter 74, which in turn alerts the thread that the data is  
20 available or has been sent. Both transfer register files  
78, 80 are connected to the EBOX 76 through a data path. In  
the described implementation, each of the register files

includes 64 registers.

The functionality of the microengine threads is determined by microcode loaded (via the core processor) for a particular user's application into each microengine's control store 70. Referring to FIG. 3, an exemplary thread task assignment 90 is shown. Typically, one of the microengine threads is assigned to serve as a receive scheduler 92 and another as a transmit scheduler 94. A plurality of threads are configured as receive processing threads 96 and transmit processing (or "fill") threads 98. Other thread task assignments include a transmit arbiter 100 and one or more core communication threads 102. Once launched, a thread performs its function independently.

The receive scheduler thread 92 assigns packets to receive processing threads 96. In a packet forwarding application for a bridge/router, for example, the receive processing thread parses packet headers and performs lookups based in the packet header information. Once the receive processing thread or threads 96 has processed the packet, it either sends the packet as an exception to be further processed by the core 23 (e.g., the forwarding information cannot be located in lookup and the core processor must

learn it), or stores the packet in the SDRAM and queues the packet in a transmit queue by placing a packet link descriptor for it in a transmit queue associated with the transmit (forwarding port) indicated by the header/lookup. The transmit queue is stored in the SRAM. The transmit arbiter thread 100 prioritizes the transmit queues and the transmit scheduler thread 94 assigns packets to transmit processing threads that send the packet out onto the forwarding port indicated by the header/lookup information during the receive processing.

The receive processing threads 96 may be dedicated to servicing particular ports or may be assigned to ports dynamically by the receive scheduler thread 92. For certain system configurations, a dedicated assignment may be desirable. For example, if the number of ports is equal to the number of receive processing threads 96, then it may be quite practical as well as efficient to assign the receive processing threads to ports in a one-to-one, dedicated assignment. In other system configurations, a dynamic assignment may provide a more efficient use of system resources.

The receive scheduler thread 92 maintains scheduling

information 104 in the GPRs 76b of the microengine within which it executes. The scheduling information 104 includes thread capabilities information 106, port-to-thread assignments (list) 108 and "thread busy" tracking information 110. At minimum, the thread capabilities information informs the receive scheduler thread as to the type of tasks for which the other threads are configured, e.g., which threads serve as receive processing threads. Additionally, it may inform the receive scheduler of other capabilities that may be appropriate to the servicing of a particular port. For instance, a receive processing thread may be configured to support a certain protocol, or a particular port or ports. A current list of the ports to which active receive processing threads have been assigned by the receive scheduler thread is maintained in the thread-to-port assignments list 108. The thread busy mask register 110 indicates which threads are actively servicing a port. The receive scheduler uses all of this scheduling information in selecting threads to be assigned to ports that require service for available packet data, as will be described in further detail below.

Referring to FIG. 4, the I/O bus interface 28

includes shared resources 120, which are coupled to a push/pull engine interface 122 and a bus interface unit 124. The bus interface unit 124 includes a ready bus controller 126 connected to a ready bus 128 and an Fbus controller 130 for connecting to a portion of the I/O bus referred to as an Fbus 132. Collectively, the ready bus 128 and the Fbus 132 make up the signals of the I/O bus 16 (FIG. 1). The resources 120 include two FIFOs, a transmit FIFO 134 and a receive FIFO 136, as well as CSRs 138, a scratchpad memory 140 and a hash unit 142. The Fbus 132 transfers data between the ports of the devices 30, 31 and the I/O bus interface 28. The ready bus 128 is an 8-bit bus that performs several functions. It is used to read control information about data availability from the devices 30, 31, e.g., in the form of ready status flags. It also provides flow control information to the devices 30, 31, and may be used to communicate with another network processor 12 that is connected to the Fbus 132. Both buses 128, 132 are accessed by the microengines 22 through the CSRs 138. The CSRs 138 are used for bus configuration, for accessing the bus interface unit 124, and for inter-thread signaling. They also include a several counters and thread status



registers, as will be described. The CSRs 138 are accessed  
by the microengines 22 and the core 23. The receive FIFO  
(RFIFO) 136 includes data buffers for holding data received  
from the Fbus 132 and is read by the microengines 22. The  
5 transmit FIFO (TFIFO) 134 includes data buffers that hold  
data to be transmitted to the Fbus 132 and is written by the  
microengines 22. The scatchpad memory 140 is accessed by  
the core 23 and microengines 22, and supports a variety of  
operations, including read and write operations, as well as  
10 bit test, bit test/clear and increment operations. The hash  
unit 142 generates hash indexes for 48-bit or 64-bit data  
and is accessed by the microengines 22 during lookup  
operations.

The processors 23 and 22 issue commands to the  
15 push/pull engine interface 122 when accessing one of the  
resources 120. The push/pull engine interface 122 places  
the commands into queues (not shown), arbitrates which  
commands to service, and moves data between the resources  
120, the core 23 and the microengines 22. In addition to  
20 servicing requests from the core 23 and microengines 22, the  
push/pull engines 122 also service requests from the ready  
bus 128 to transfer control information to a register in the

microengine read transfer registers 80.

When a thread issues a request to a resource 120, a command is driven onto an internal command bus 150 and placed in queues within the push/pull engine interface 122. Receive/read-related instructions (such as instructions for reading the CSRS) are written to a "push" command queue.

The CSRs 138 include the following types of registers: Fbus receive and transmit registers; Fbus and ready bus configuration registers; ready bus control registers; hash unit configuration registers; interrupt registers; and several miscellaneous registers, including a thread status registers. Those of the registers which pertain to the receive process will be described in further detail.

The interrupt/signal registers include an INTER\_THD\_SIG register for inter-thread signaling. Any thread within the microengines 22 or the core 23 can write a thread number to this register to signal an inter-thread event.

Further details of the Fbus controller 130 and the ready bus controller 126 are shown in FIG. 5. The ready bus controller 126 includes a programmable sequencer 160 for

retrieving MAC device status information from the MAC  
devices 30, 31, and asserting flow control to the MAC  
devices over the ready bus 128 via ready bus interface logic  
161. The Fbus controller 130 includes Fbus interface logic  
162, which is used to transfer data to and from the devices  
30, 31, is controlled by a transmit state machine (TSM) 164  
and a receive state machine (RSM) 166. In the embodiment  
herein, the Fbus 132 may be configured as a bidirectional  
64-bit bus, or two dedicated 32-bit buses. In the  
unidirectional, 32-bit configuration, each of the state  
machines owns its own 32-bit bus. In the bidirectional  
configuration, the ownership of the bus is established  
through arbitration. Accordingly, the Fbus controller 130  
further includes a bus arbiter 168 for selecting which state  
machine owns the Fbus 132.

Some of the relevant CSRs used to program and  
control the ready bus 128 and Fbus 132 for receive processes  
are shown in FIGS. 6A-6F. Referring to FIG. 6A,  
RDYBUS\_TEMPLATE\_PROGx registers 170 are used to store  
instructions for the ready bus sequencer. Each register of  
these 32-bit registers 170a, 170b, 170c, includes four, 8-  
bit instruction fields 172. Referring to FIG. 6B, a

RCV\_RDY\_CTL register 174 specifies the behavior of the receive state machine 166. The format is as follows: a reserved field (bits 31:15) 174a; a fast port mode field (bits 14:13) 174b, which specifies the fast (Gigabit) port thread mode, as will be described; an auto push prevent window field (bits 12:10) 174c for specifying the autopush prevent window used by the ready bus sequencer to prevent the receive scheduler from accessing its read transfer registers when an autopush operation (which pushes information to those registers) is about to begin; an autopush enable (bit 9) 174d, used to enable autopush of the receive ready flags; another reserved field (bit 8) 174e; an autopush destination field (bits 7:6) 174f for specifying an autopush operation's destination register; a signal thread enable field (bit 5) 174g which, when set, indicates the thread to be signaled after an autopush operation; and a receive scheduler thread ID (bits 4:0) 174h, which specifies the ID of the microengine thread that has been configured as a receive scheduler.

Referring to FIG. 6C, a REC\_FASTPORT\_CTL register 176 is relevant to receiving packet data from fast ports (fast port mode) only. It enables receive threads to view

the current assignment of header and body thread assignments for the two fast ports, as will be described. It includes the following fields: a reserved field (bits 31:20) 176a; an FP2\_HDR\_THD\_ID field (bits 19:15) 176b, which specifies the fast port 2 header receive (processing) thread ID; an FP2\_BODY\_THD\_ID field (bits 14:10) 176c for specifying the fast port 2 body receive processing thread ID; an FP1\_HDR\_THD\_ID field (bits 9:5) 176d for specifying the fast port 1 header receive processing thread ID; and an FP1\_BODY\_THD\_ID field (bits 4:0) 176e for specifying the fast port 1 body processing thread ID. The manner in which these fields are used by the RSM 166 will be described in detail later.

Although not depicted in detail, other bus registers include the following: a RDYBUS\_TEMPLATE\_CTL register 178 (FIG. 6D), which maintains the control information for the ready bus and the Fbus controllers, for example, it enables the ready bus sequencer; a RDYBUS\_SYNCH\_COUNT\_DEFAULT register 180 (FIG. 6E), which specifies the program cycle rate of the ready bus sequencer; and an FP\_FASTPORT\_CTL register 182 (FIG. 6F), which specifies how many Fbus clock cycles the RSM 166 must wait between the last data transfer

and the next sampling of fast receive status, as will be described.

Referring to FIG. 7A, the MAC device 30 provides transmit status flags 200 and receive status flags 202 that indicate whether the amount of data in an associated transmit FIFO 204 or receive FIFO 206 has reached a certain threshold level. The ready bus sequencer 160 periodically polls the ready flags (after selecting either the receive ready flags 202 or the transmit ready flags 200 via a flag select 208) and places them into appropriate ones of the CSRs 138 by transferring the flag data over ready bus data lines 209. In this embodiment, the ready bus includes 8 data lines for transferring flag data from each port to the Fbus interface unit 124. The CSRs in which the flag data are written are defined as RCV\_RDY\_HI/LO registers 210 for receive ready flags and XMIT\_RDY\_HI/LO registers 212 for transmit ready flags, if the ready bus sequencer 160 is programmed to execute receive and transmit ready flag read instructions, respectively.

When the ready bus sequencer is programmed with an appropriate instruction directing it to interrogate MAC receive ready flags, it reads the receive ready flags from

the MAC device or devices specified in the instruction and places the flags into RCV\_RDY\_HI register 210a and a RCV\_RDY\_LO register 210b, collectively, RCV\_RDY registers 210. Each bit in these registers corresponds to a different device port on the I/O bus.

Also, and as shown in FIG. 7B, the bus interface unit 124 also supports two fast port receive ready flag pins FAST\_RX1 214a and FAST\_RX2 214b for the two fast ports of the fast MAC device 31. These fast port receive ready flag pins are read by the RSM 166 directly and placed into an RCV\_RDY\_CNT register 216.

The RCV\_RDY\_CNT register 216 is one of several used by the receive scheduler to determine how to issue a receive request. It also indicates whether a flow control request is issued.

Referring to FIG. 8A, the format of the RCV\_RDY\_CNT register 216 is as follows: bits 31:28 are defined as a reserved field 216a; bit 27 is defined as a ready bus master field 216b and is used to indicate whether the ready bus 128 is configured as a master or slave; a field corresponding to bit 26 216c provides flow control information; bits 25 and 24 correspond to FRDY2 field 216d and FRDY1 field 216e,

respectively. The FRDY2 216d and FRDY1 216e are used to store the values of the FAST\_RX2 pin 214b and FAST\_RX1 pin 214a, respectively, both of which are sampled by the RSM 166 each Fbus clock cycle; bits 23:16 correspond to a reserved field 216f; a receive request count field (bits 15:8) 216g specifies a receive request count, which is incremented after the RSM 166 completes a receive request and data is available in the RFIFO 136; a receive ready count field (bits 7:0) 216h specifies a receive ready count, an 8-bit counter that is incremented each time the ready bus sequencer 160 writes the ready bus registers RCV\_RDY\_CNT register 216, the RCV\_RDY\_LO register 210b and RCV\_RDY\_HI register 210a to the receive scheduler read transfer registers.

There are two techniques for reading the ready bus registers: "autopush" and polling. The autopush instruction may be executed by the ready bus sequencer 160 during a receive process (rxautopush) or a transmit process (txautopush). Polling requires that a microengine thread periodically issue read references to the I/O bus interface 28.

The rxautopush operation performs several functions.



It increments the receive ready count in the RCV\_RDY\_CNT register 216. If enabled by the RCV\_RDY\_CTL register 174, it automatically writes the RCV\_RDY\_CNT 216, the RCV\_RDY\_LO and RCV\_RDY\_HI registers 210b, 210a to the receive scheduler read transfer registers and signals to the receive scheduler thread 92 (via a context event signal) when the rxautopush operation is complete.

The ready bus sequencer 160 polls the MAC FIFO status flags periodically and asynchronously to other events occurring in the processor 12. Ideally, the rate at which the MAC FIFO ready flags are polled is greater than the maximum rate at which the data is arriving at the MAC ports. Thus, it is necessary for the receive scheduler thread 92 to determine whether the MAC FIFO ready flags read by the ready bus sequencer 160 are new, or whether they have been read already. The rxautopush instruction increments the receive ready count in the RCV\_RDY\_CNT register 216 each time the instruction executes. The RCV\_RDY\_CNT register 216 can be used by the receive scheduler thread 92 to determine whether the state of specific flags have to be evaluated or whether they can be ignored because receive requests have been issued and the port is currently being serviced. For

example, if the FIFO threshold for a Gigabit Ethernet port is set so that the receive ready flags are asserted when 64 bytes of data are in the MAC receive FIFO 206, then the state of the flags does not change until the next 64 bytes arrive 5120 ns later. If the ready bus sequencer 160 is programmed to collect the flags four times each 5120 ns period, the next three sets of ready flags that are to be collected by the ready bus sequence 160 can be ignored.

When the receive ready count is used to monitor the freshness of the receive ready flags, there is a possibility that the receive ready flags will be ignored when they are providing new status. For a more accurate determination of ready flag freshness, the receive request count may be used. Each time a receive request is completed and the receive control information is pushed onto the RCV\_CNTL register 232, the the RSM 166 increments the receive request count. The count is recorded in the RCV\_RDY\_CNT register the first time the ready bus sequencer executes an rxrdy instruction for each program loop. The receive scheduler thread 92 can use this count to track how many requests the receive state machine has completed. As the receive scheduler thread issues commands, it can maintain a list of the receive

requests it submits and the ports associated with each such request.

Referring to FIGS. 8B and 8C, the registers RCV\_RDY\_HI 210a and RCV\_RDY\_LO 210b have a flag bit 217a, 217b, respectively, corresponding to each port.

Referring to FIG. 9, the receive scheduler thread 92 performs its tasks as quickly as possible to ensure that the RSM 166 is always busy, that is, that there is always a receive request waiting to be processed by the RSM 166. Several tasks performed by the receive scheduler 92 are as follows. The receive scheduler 92 determines which ports need to be serviced by reading the RCV\_RDY\_HI, RCV\_RDY\_LO and RCV\_RDY\_CNT registers 210a, 210b and 216, respectively. The receive scheduler 92 also determines which receive ready flags are new and which are old using either the receive request count or the receive ready count in the RCV\_RDY\_CNT register, as described above. It tracks the thread processing status of the other microengine threads by reading thread done status CSRs 240. The receive scheduler thread 92 initiates transfers across the Fbus 132 via the ready bus, while the receive state machine 166 performs the actual read transfer on the Fbus 132. The receive scheduler

92 interfaces to the receive state machine 166 through two  
FBI CSRs 138: an RCV\_REQ register 230 and an RCV\_CNTL  
register 232. The RCV\_REQ register 230 instructs the  
receive state machine on how to receive data from the Fbus  
132.

Still referring to FIG. 9, a process of initiating  
an Fbus receive transfer is shown. Having received ready  
status information from the RCV\_RDY\_HI/LO registers 210a,  
210b as well as thread availability from the thread done  
register 240 (transaction "1", as indicated by the arrow  
labeled 1), the receive scheduler thread 92 determines if  
there is room in the RCV\_REQ FIFO 230 for another receive  
request. If it determines that RCV\_REQ FIFO 230 has room to  
receive a request, the receive scheduler thread 92 writes a  
receive request by pushing data into the RCV\_REQ FIFO 230  
(transaction 2). The RSM 166 processes the request in the  
RCV\_REQ FIFO 230 (transaction 3). The RSM 166 responds to  
the request by moving the requested data into the RFIFO 136  
(transaction 4), writing associated control information to  
the RCV\_CTL FIFO 232 (transaction 5) and generating a  
start\_receive signal event to the receive processing thread  
96 specified in the receive request (transaction 6). The

RFIFO 136 includes 16 elements 241, each element for storing  
a 64 byte segment of data referred to herein as a MAC packet  
("MPKT"). The RSM 166 reads packets from the MAC ports in  
fragments equal in size to one or two RFIFO elements, that  
is, MPKTs. The specified receive processing thread 96  
responds to the signal event by reading the control  
information from the RCV\_CTL register 232 (transaction 7).  
It uses the control information to determine, among other  
pieces of information, where the data is located in the  
RFIFO 136. The receive processing thread 96 reads the data  
from the RFIFO 136 on quadword boundaries into its read  
transfer registers or moves the data directly into the SDRAM  
(transaction 8).

The RCV\_REQ register 230 is used to initiate a  
receive transfer on the Fbus and is mapped to a two-entry  
FIFO that is written by the microengines. The I/O bus  
interface provides signals (not shown) to the receive  
scheduler thread indicating that the RCV\_REQ FIFO 236 has  
room available for another receive request and that the last  
issued receive request has been stored in the RCV\_REQ  
register 230.

Referring to FIG. 10A, the RCV\_REQ FIFO 230 includes

two entries 231. The format of each entry 231 is as follows. The first two bits correspond to a reserved field 230a. Bit 29 is an FA field 230b for specifying the maximum number of Fbus accesses to be performed for this request. A THSG field (bits 28:27) 230c is a two-bit thread message field that allows the scheduler thread to pass a message to the assigned receive thread through the ready state machine, which copies this message to the RCV\_CNTL register. An SL field 230d (bit 26) is used in cases where status information is transferred following the EOP MPKT. It indicates whether two or one 32-bit bus accesses are required in a 32-bit Fbus configuration. An E1 field 230e (bits 21:18) and an E2 field (bits 25:22) 230f specify the RFIFO element to receive the transferred data. If only 1 MPKT is received, it is placed in the element indicated by the E1 field. If two MPKTs are received, then the second MPKT is placed in the RFIFO element indicated by the E2 field. An FS field (bits 17:16) 230g specifies use of a fast or slow port mode, that is, whether the request is directed to a fast or slow port. The fast port mode setting signifies to the RSM that a sequence number is to be associated with the request and that it will be handling

speculative requests, which will be discussed in further detail later. An NFE field (bit 15) 230h specifies the number of RFIFO elements to be filled (i.e., one or two elements). The IGFR field (bit 13) 230i is used only if fast port mode is selected and indicates to the RSM that it should process the request regardless of the status of the fast ready flag pins. An SIGRS field (bit 11) 230j, if set, indicates that the receive scheduler be signaled upon completion of the receive request. A TID field (bits 10:6) 230k specifies the receive thread to be notified or signaled after the receive request is processed. Therefore, if bit 11 is set, the RCV\_REQ entry must be read twice, once by the receive thread and once by the receive scheduler thread, before it can be removed from the RCV\_REQ FIFO. An RM field (bits 5:3) 230l specified the ID of the MAC device that has been selected by the receive scheduler. Lastly, an RP field (bits 2:0) 230m specifies which port of the MAC device specified in the RM field 230l has been selected.

The RSM 166 reads the RCV\_REQ register entry 231 to determine how it should receive data from the Fbus 132, that is, how the signaling should be performed on the Fbus, where the data should be placed in the RFIFO and which microengine

thread should be signaled once the data is received. The RSM 166 looks for a valid receive request in the RCV\_REQ FIFO 230. It selects the MAC device identified in the RM field and selects the specified port within the MAC by asserting the appropriate control signals. It then begins receiving data from the MAC device on the Fbus data lines. The receive state machine always attempts to read either eight or nine quadwords of data from the MAC device on the Fbus as specified in the receive request. If the MAC device asserts the EOP signal, the RSM 166 terminates the receive early (before eight or nine accesses are made). The RSM 166 calculates the total bytes received for each receive request and reports the value in the REC\_CNTL register 232. If EOP is received, the RSM 166 determines the number of valid bytes in the last received data cycle.

The RCV\_CNTL register 232 is mapped to a four-entry FIFO (referred to herein as RCV\_CNTL\_FIFO 232) that is written by the receive state machine and read by the microengine thread. The I/O bus interface 28 signals the assigned thread when a valid entry reaches the top of the RCV\_CNTL\_FIFO. When a microengine thread reads the RCV\_CNTL register, the data is popped off the FIFO. If the SIGRS



field 230i is set in the RCV\_REQ register 230, the receive scheduler thread 92 specified in the RCV\_CNTL register 232 is signaled in addition to the thread specified in TID field 230k. In this case, the data in the RCV\_CNTL register 232 is read twice before the receive request data is retired from the RCV\_CTL FIFO 232 and the next thread is signaled. The receive state machine writes to the RCV\_CTL register 232 as long as the FIFO is not full. If the RCV\_CTL FIFO 232 is full, the receive state machine stalls and stops accepting any more receive requests.

Referring to FIG. 10B, the RCV\_CNTL FIFO 232 provides instruction to the signaled thread (i.e., the thread specified in TID) to process the data. As indicated above, the RCV\_CNTL FIFO includes 4 entries 233. The format of the RCV\_CNTL FIFO entry 233 is as follows: a THMSG field (31:30) 23a includes the 2-bit message copied by the RSM from REC\_REQ register[28:27]. A MACPORT/THD field (bits 29:24) 232b specifies either the MAC port number or a receive thread ID, as will be described in further detail below. An SOP SEQ field (23:20) 232c is used for fast ports and indicates a packet sequence number as an SOP (start-of-packet) sequence number if the SOP was asserted during the

receive data transfer and indicates an MPKT sequence number  
if SOP was not so asserted. An RF field 232d and RERR field  
232e (bits 19 and 18, respectively) both convey receive  
error information. An SE field 232f (17:14) and an FE field  
5 232g (13:10) are copies of the E2 and E1 fields,  
respectively, of the REC\_REQ. An EF field (bit 9) 232h  
specifies the number of RFIFO elements which were filled by  
the receive request. An SN field (bit 8) 232i is used for  
fast ports and indicates whether the sequence number  
10 specified in SOP\_SEQ field 232c is associated with fast port  
1 or fast port 2. A VLD BYTES field (7:2) 232j specifies  
the number of valid bytes in the RFIFO element if the  
element contains in EOP MPKT. An EOP field (bit 1) 232k  
indicates that the MPKT is an EOP MPKT. An SOP field (bit  
15 0) 232l indicates that the MPKT is an SOP MPKT.

FIG. 11 illustrates the format of the thread done  
registers 240 and their interaction with the receive  
scheduler and processing threads 92, 96, respectively, of  
the microengines 22. The thread done registers 240 include  
20 a first thread status register, TH\_DONE\_REG0 240a, which has  
2-bit status fields 241a corresponding to each of threads 0  
through 15. A second thread status register, TH\_DONE\_REG1

240b, has 2-bit status fields 241b corresponding to each of  
threads 16 through 23. These registers can be read and  
written to by the threads using a CSR instruction (or fast  
write instruction, described below). The receive scheduler  
5 thread can use these registers to determine which RFIFO  
elements are not in use. Since it is the receive scheduler  
thread 92 that assigns receive processing threads 96 to  
process the data in the RFIFO elements, and it also knows  
the thread processing status from the THREAD\_DONE\_REG0 and  
10 THREAD\_DONE\_REG1 registers 240a, 240b, it can determine  
which RFIFO elements are currently available.

The THREAD\_DONE CSRs 240 support a two-bit message  
for each microengine thread. The assigned receive thread  
may write a two-bit message to this register to indicate  
15 that it has completed its task. Each time a message is  
written to the THREAD\_DONE register, the current message is  
logically ORed with the new message. The bit values in the  
THREAD\_DONE registers are cleared by writing a "1", so the  
scheduler may clear the messages by writing the data read  
20 back to the THREAD\_DONE register. The definition of the 2-  
bit status field is determined in software. An example of  
four message types is illustrated in TABLE 1 below.

2-BIT MESSAGE	DEFINITION
00	Busy.
01	Idle, processing complete.
10	Not busy, but waiting to finish processing of entire packet.
11	Idle, processing complete for an EOP MPKT.

TABLE 1

The assigned receive processing threads write their status to the THREAD\_DONE register whenever the status changes. For example, a thread may immediately write 00 to the THREAD\_DONE register after the receive state machine signals the assigned thread. When the receive scheduler thread reads the THREAD\_DONE register, it can look at the returned value to determine the status of each thread and then update its thread/port assignment list.

The microengine supports a fast\_wr instruction that improves performance when writing to a subset of CSR registers. The fast\_wr instruction does not use the push or pull engines. Rather, it uses logic that services the instruction as soon as the write request is issued to the FBI CSR. The instruction thus eliminates the need for the

pull engine to read data from a microengine transfer register when it processes the command. The meaning of the 10-bit immediate data for some of the CSRs is shown below.

CSR	10-BIT IMMEDIATE DATA
INTER_THD_SIG	Thread number of the thread that is to be signaled.
THREAD_DONE	A 2-bit message that is shifted into a position relative to the thread that is writing the message.
THREAD_DONE_INCR1 THREAD_DONE_INCR2	Same as THREAD_DONE except that either the enqueue_seq1 or enqueue_seq2 is also incremented.
INCR_ENQ_NUM1 INCR_ENQ_NUM2	Write a one to increment the enqueue sequence number by one.

TABLE 2

It will be appreciated that the receive process as described herein assumes that no packet exemptions occurred, that is, that the threads are able to handle the packet processing without assistance from the core processor. Further, the receive process as described also assumes the availability of FIFO space. It will be appreciated that the various state machines must determine if there is room

available in a FIFO, e.g., the RFIFO, prior to writing new entries to that FIFO. If a particular FIFO is full, the state machine will wait until the appropriate number of entries has been retired from that FIFO.

5           Additions, subtractions, and other modifications of the preferred embodiments of the invention will be apparent to those practiced in this field and are within the scope of the following claims.

What is claimed is:

1           1.    A method of receiving data from a network,  
2    comprising:  
3           issuing a receive request directing a transfer of data  
4    from one of a plurality of device ports to a buffer memory and  
5    specifying a thread from among a plurality of processing program  
6    threads to process the data.

1           2.    The method of claim 1, further comprising:  
2           determining if any of the plurality of device ports  
3           coupled to the network require service.

1           3.    The method of claim 2, further comprising:  
2           transferring the data to the buffer memory and signaling  
3           to the specified program thread that the data is ready for  
4           processing.

1           4.    The method of claim 2, wherein determining  
2           comprises:  
3           interrogating the plurality of device ports to identify  
4           which of the plurality of device ports require service.

1           5.    The method of claim 4, wherein determining further  
2 comprises:

3           preparing control information corresponding to those  
4 device ports identified as requiring service.

1           6.    The method of claim 5, wherein the control  
2 information comprises receive ready flags each associated with a  
3 device port receive FIFO in a corresponding one of the device  
4 ports.

1           7.    The method of claim 6, wherein interrogating  
2 comprises:

3           polling the state of the ready flags to determine if the  
4 ready flags are asserted, the assertion of the ready flags  
5 indicating that the corresponding device ports have data ready  
1 for transfer.

1           8.    The method of claim 7, wherein the receive ready  
2 flags indicate that the associated device port receive FIFO has  
3 reached a threshold level of fullness.



1           9.    The method of claim 8, wherein the receive ready  
2 flags indicate that the associated device port receive FIFO  
3 stores a full network packet.

1           10. The method of claim 5, further comprising:  
2           maintaining a receive ready count, the receive ready  
3 count being incremented when the control information is prepared.

1           11. The method of claim 5, wherein preparing control  
2 information further comprises:  
3           writing a flag to a control and status register for each  
4 device port in the plurality of device ports that is determined  
5 to require service.

1           12. The method of claim 11, wherein issuing comprises:  
2           obtaining the control information from the control and  
3 status register; and  
4           selecting from each device port in the plurality of  
5 device ports having set bits in the control and status register a  
6 port for servicing.

1           13. The method of claim 12, wherein issuing further

comprises:

determining which among the plurality of program threads  
is available; and  
assigning an available program thread to the selected  
port.

14. The method of claim 12, wherein selecting a port  
comprises:

using the receive ready count to determine if the ready  
flags reflect current status of the device port.

15. The method of claim 3, further comprising:  
maintaining a receive request count for counting transfer  
of data to the buffer memory, the receive request count being  
incremented by one upon the transfer of the data to the buffer  
memory and signaling to the specified program thread.

16. The method of claim 15, wherein selecting a port  
further comprises:

using the receive request count to determine if the ready  
flags reflect current status of the device ports.

1           17. A method of receiving data from a plurality of  
2 peripheral ports, comprising:

3           determining that the one of the plurality of peripheral  
4 ports requires servicing;

5           issuing a receive request based on the determination, the  
6 receive request directing the transfer of data from the one of  
7 the plurality of peripheral ports to a buffer memory and  
8 specifying a program thread from among of a plurality of  
9 processing program threads to process the data; and

10          transferring the data to the buffer memory and signaling  
11 to the specified thread that the data is ready for processing.

12          18. An article comprising a computer-readable medium  
13 which stores computer-executable instructions for receiving data  
14 from a plurality of ports, the instructions causing a computer  
15 to:

16          issue a receive request directing a transfer of data from  
17 one of a plurality of device ports to a buffer memory and  
18 specifying a program thread from among a plurality of processing  
19 program threads to process the data.

20          19. The article of claim 18, the article further

2 comprises instructions causing a computer to:

3 determine if any of the plurality of device ports coupled  
4 to the network require service.

1 20. The article of claim 19, the article further  
2 comprises instructions causing a computer to:

3 transfer the data to the buffer memory and signal to the  
4 specified program thread that the data is ready for processing.

5 21. The article of claim 19, wherein the instructions to  
6 determine comprise instructions causing a computer to:

7 interrogate the plurality of device ports to identify  
8 which of the plurality of device ports require service; and

9 prepare control information corresponding to those device  
10 ports identified as requiring service.

1 22. The article of claim 21, the article further  
2 comprises instructions causing a computer to:

3 maintain a receive ready count, the receive ready count  
4 being incremented when the control information is prepared.

1 23. The article of claim 22, wherein the instructions to

2 issue comprise instructions causing a computer to:

3 use the receive ready count to check the current status  
4 of the device port.

1 24. The article of 19, the article further comprises  
2 instructions causing a computer to:

3 maintain a receive request count for counting transfer of  
4 data to the buffer memory, the receive request count being  
5 incremented by one upon the transfer of the data to the buffer  
6 memory and signaling to the specified program thread.

1 25. The article of claim 24, wherein the instructions to  
2 issue comprise instructions causing a computer to:

3 use the receive request count to check the current status  
4 of the device ports.

1 26. A network processor comprising:

2 a microengine for executing threads, the threads  
3 including a receive scheduler program thread and receive  
4 processing program threads; and

5 the receive scheduler thread assigning a port to one of  
6 the receive processing program threads if the port has available


7 data.

1 27. The network processor, further wherein the receive  
2 scheduler program thread directs transfer of the data  
3 to the assigned one of the receive processing program threads for  
4 processing.

1 28. The network processor, further comprising:  
2 an interface coupled to the microengine for receiving  
3 data from the port, the interface for indicating to the receive  
4 scheduler program thread whether the port has data available for  
5 processing by one of the receive processing program threads.

Abstract of the Disclosure

A network processor that has multiple processing elements, each supporting multiple simultaneous program threads with access to shared resources in an interface. Control logic in the interface samples the ready state of network ports and forwards the ready state information to a scheduler program thread. The scheduler program thread issues receive request command that direct the interface to fetch segments of data from a selected one of the network ports into a receive FIFO for processing by an assigned one of a plurality of receive processing program threads. For each segment of data that is to be transferred to the receive FIFO, a control FIFO is loaded with control information specifying the associated receive FIFO location(s), the selected one of the network ports and the assigned receive processing program thread. The request is processed by reading the control information, transferring the data indicated by the request to the receive FIFO and loading another control FIFO with control information specifying how the data is to be processed by the assigned receive processing program thread. The interface signals the assigned receive processing program thread to process the data.





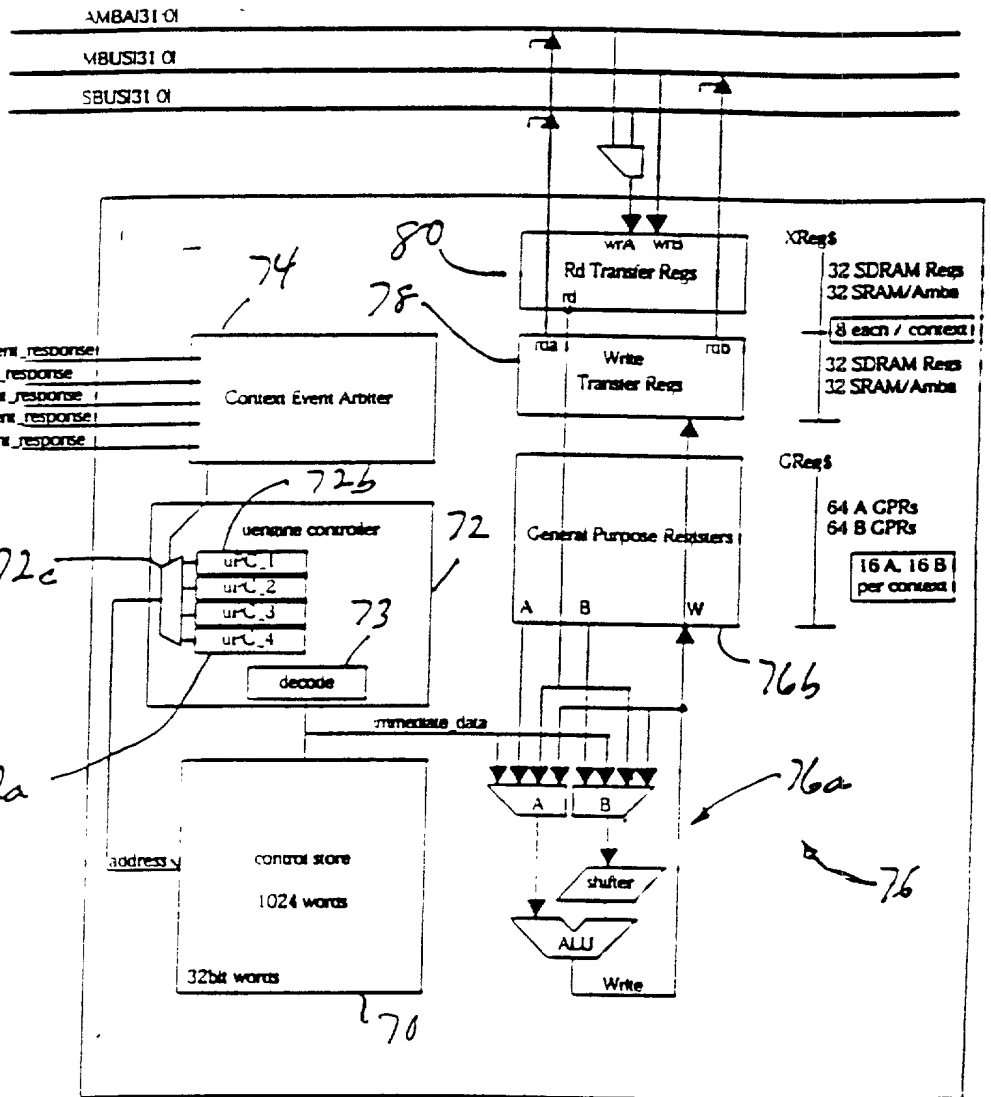


FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6A




FIG. 6B




FIG. 6C




FIG. 6D




FIG. 6E




FIG. 6F




FIG. 7A




FIG. 7B






FIG. 8A




FIG. 8B




FIG. 8C




FIG. 9

231	230a	230b	230c	230d	230e	230f	230g	230h	230i	230j	230k	230l	230m	230
RES	FA		DL	EZ	EI	FS	NFE	IGFA	STCAS	TID	RM	RP		0
														1

FIG. 10A

233	232a	232b	232c	232d	232e	232f	232g	232h	232i	232j	232k	232l	232
THMSG	MACPORT TYPE	SOP SEQ#	NF	NEAR	SE	FE	EFOR	SNOR	Valid bytes	EOP	SOP		0
													1
													2
													3

FIG. 10B

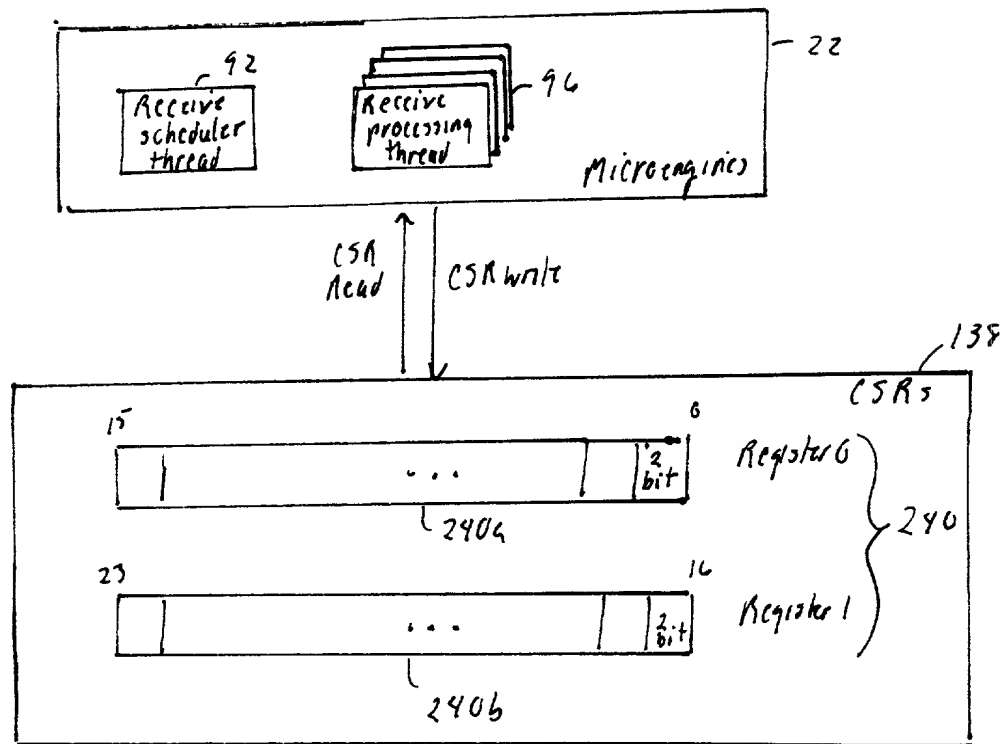


FIG. 11